REMARKS/ARGUMENTS

In the Office action dated September 9, 2005, claims 1 – 13 and 15 – 17 were rejected. In response, claims 15 – 17 have been renumbered as claims 14 – 16, respectively, claims 9, 13, 14, 15, and 16 have been amended, claim 12 has been cancelled, and new claims 17 – 26 have been added. Applicant hereby requests reconsideration of the application in view of the amended claims, the added claims, and the below-provided remarks.

I. Amendments to the Specification

Applicant notes with appreciation the identification of errors in the specification as filed. The specification has been amended to correct the identified errors. In particular, reference to FIG. 4 has been changed to FIG. 3 and reference to task 503 has been changed to task 502.

II. Amendments to the Claims

Claim 9 has been amended to recite "said first *read* pointer" instead of "said first *write* pointer."

Claim 13 has been amended to clean up inconsistent use of the terms "first" and "second" and to depend from claim 11 instead of claim 12.

Claim 14 has been renumbered (from claim 15) and amended to recite "said first *read* pointer" instead of "said first *write* pointer."

Claim 15 has been renumbered (from claim 16) and amended to recite "said first *read* pointer" instead of "said first *write* pointer" and to recite "said third *read* pointer" instead of "said third *write* pointer."

Claim 16 has been renumbered (from claim 17) and amended to depend from claim 15 instead of claim 16.

III. Claim Rejections Under 35 U.S.C. 102

Claims 1 – 13 and 15 – 17 (as filed) were rejected under 35 U.S.C. 102(b) as being anticipated by Kusyk (U.S. Pat. No. 6,246,668).

Claim 1

Claim 1 recites:

> "An integrated circuit comprising:
> a first input port for receiving a first time-division multiplexed signal that comprises a first series of frame boundaries;
> a second input port for receiving a second time-division multiplexed signal that comprises a second series of frame boundaries;
> *a first frame position register* whose contents are related to how far *said first time-division multiplexed signal* is from a frame boundary in said first time-division multiplexed signal *at a point in time*; and
> *a second frame position register* whose contents are related to how far *said second time-division multiplexed signal* is from a frame boundary in said second time-division multiplexed signal *at said point in time*."
> (emphasis added)

Applicant asserts that claim 1 is not anticipated by Kusyk because Kusyk does not disclose "a first frame position register whose contents are related to…said first time-division multiplexed signal" and "a second frame position register whose contents are related to…said second time-division multiplexed signal." As recited in claim 1, the contents of the first and second frame position registers represent measures related to two <u>different</u> time-division multiplexed signals, i.e., the first time-division multiplexed signal and the second time-division multiplexed signal.

With regard to claim 1, the Office action cites that the first frame register recited in claim 1 is disclosed by "H1" in Kusyk and that the second frame register recited in claim 1 is disclosed by "H2" and column 7, lines 62 – 64 and column 5, lines 30 – 32 in Kusyk.

As is well-known in the field, the SONET telecommunications standard uses a payload pointer to compensate for frequency and phase variations that can occur between network nodes that use separate clocks. The payload pointer is carried in the H1 and H2 bytes of the transport overhead of an STS frame and allows a synchronous payload envelope (SPE) to be separated from the transport overhead by a varying

number of bytes. Use of the payload pointer carried in the H1 and H2 bytes of the transport overhead of an STS frame is briefly described by Kusyk at column 3, lines 32 – 41 and is described in more detail in the document entitled "SONET Telecommunications Standard, Primer," (see pages 14 – 15) which is included, in part, as Appendix A.

While Kusyk does disclose the H1 and H2 bytes of an STS frame, Applicant asserts that the H1 and H2 bytes of an STS frame do not disclose "a first frame position register whose contents are related to...*said first time-division multiplexed signal*" and "a second frame position register whose contents are related to...*said second time-division multiplexed signal*" as recited in claim 1. The first frame position register recited in claim 1 is for information related to a first time-division multiplexed signal while the second frame position register recited in claim 1 is for information related to a second time-division multiplexed signal. In contrast to claim 1, the H1 and H2 bytes of an STS frame relate to a single SPE. Because the first and second frame position registers recited in claim 1 relate to respective first and second time-division multiplexed signals while the H1 and H2 bytes of an STS frame relate to a single SPE, Applicant asserts that claim 1 is not anticipated by Kusyk.

Further, Kusyk discloses the concept of transmitting two trace messages, one trace message that travels a short path and one trace path message that travels a long path, as a way to determine and compensate for the time delay between two different paths that have the same source and destination. The delay between the trace messages is determined, as described in column 10, line 34 to column 11, line 12, by comparing read and write addresses. The delay between the trace messages is compensated for by adjusting the payload pointer (H1 and H2 bytes) as described in column 8, lines 7 – 21. Although Kusyk discloses adjusting the payload pointer (H1 and H2 bytes) to compensate for the delay between the trace messages, nowhere does Kusyk disclose "a first frame position register..." and "a second frame position register..." whose contents are measures taken at the same point in time as recited in claim 1.

IV. Independent Claims 5, 9, 11, and 15

Independent claims 5, 8, 11, and 15 include similar limitations to claim 1. Because of the similarities between claim 1 and claims 5, 8, 11, and 15, the remarks provided above with reference to claim 1 apply also to claims 5, 8, 11, and 15.

V. New Claims 17 - 26

Claim 17

Claim 17 recites a controller that is configured to synchronize the first and second time-division multiplexed signals. Kusyk does not disclose a single controller that is configured to perform synchronization of two different time-division multiplexed signals that travel two different paths. In fact, Kusyk discloses that a single controller is not desirable. In particular, at column 2, lines 20 – 24, Kusyk discloses:

> "The method and apparatus according to the present invention advantageously use linked point processors, which eliminate the need for a master-slave relationship. A master-slave relationship is generally more difficult to control."

That is, Kusyk teaches away from a common synchronization controller, while claim 17 recites a common synchronization controller.

Claim 18

Applicant asserts that Kusyk does not disclose the generation of offset values as recited in claim 18.

Claim 19

Applicant asserts that Kusyk does not disclose the adjustment of read pointers as recited in claim 19.

Claims 20 – 26

New claims 20 – 26 include limitations similar to claims 17 – 19 and therefore the above-provided remarks apply also to these claims.
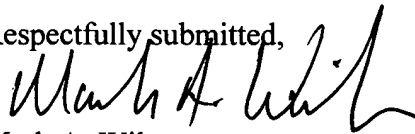
V. Dependent Claims 2 – 4, 6, 7, 9, 10, 13, 14, and 16

Applicant asserts that dependent claims 2 – 4, 6, 7, 9, 10, 13, 14, and 16 are allowable at least based on their respective allowable base claims.

Applicant respectfully requests reconsideration of the claims in view of the remarks made herein. A notice of allowance is earnestly solicited.

At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account **50-3444** pursuant to 37 CFR 1.25. Additionally, please charge any fees to Deposit Account **50-3444** under 37 CFR 1.16, 1.17, 1.19, 1.20 and 1.21.

Respectfully submitted,

Date: January 9, 2006

Mark A. Wilson
Reg. No. 43,994

Wilson & Ham
PMB: 348
2530 Berryessa Road
San Jose, CA 95132
Phone: (925) 249-1300
Fax: (925) 249-0111

APPENDIX A

for Application Serial Number 09/941,894

"SONET Telecommunications Standard Primer," Tektronix 2001, pages 1 – 19

SONET Telecommunications Standard Primer

Tektronix

What is SONET?

This document provides an introduction to the Synchronous Optical NETwork (SONET) standard. Standards in the telecommunications field are always evolving. Information in this SONET primer is based on the latest information available from the Bellcore and ITU-T standards organizations.

Use this primer as an introduction to the technology of SONET. If more detailed information is required, consult the latest material from Bellcore and ITU-T, paying particular attention to the latest date.

For help in understanding the language of SONET telecommunications, a comprehensive Glossary is provided at the end of this document.

SONET Telecommunications Standard
► Primer

Contents

Introduction To SONET

SONET (Synchronous Optical NETwork) is a standard for optical telecommunications transport. It was formulated by the Exchange Carriers Standards Association (ECSA) for the American National Standards Institute (ANSI), which sets industry standards in the U.S. for telecommunications and other industries. The comprehensive SONET/SDH standard is expected to provide the transport infrastructure for worldwide telecommunications for at least the next two or three decades.

The increased configuration flexibility and bandwidth availability of SONET provides significant advantages over the older telecommunications system. These advantages include:

▶ Reduction in equipment requirements and an increase in network reliability

▶ Provision of overhead and payload bytes – the overhead bytes permit management of the payload bytes on an individual basis and facilitate centralized fault sectionalization

▶ Definition of a synchronous multiplexing format for carrying lower level digital signals (such as DS1, DS3) and a synchronous structure which greatly simplifies the interface to digital switches, digital cross-connect switches, and add-drop multiplexers

▶ Availability of a set of generic standards which enable products from different vendors to be connected

▶ Definition of a flexible architecture capable of accommodating future applications, with a variety of transmission rates

In brief, SONET defines optical carrier (OC) levels and electrically equivalent synchronous transport signals (STSs) for the fiber-optic based transmission hierarchy.

Background

Before SONET, the first generations of fiber optic systems in the public telephone network used proprietary architectures, equipment, line codes, multiplexing formats, and maintenance procedures. The users of this equipment – Regional Bell Operating Companies and inter-exchange carriers (IXCs) in the U.S., Canada, Korea, Taiwan, and Hong Kong – wanted standards so they could mix and match equipment from different suppliers. The task of creating such a standard was taken up in 1984 by the Exchange Carriers Standards Association (ECSA) to establish a standard for connecting one fiber system to another. This standard is called SONET for Synchronous Optical NETwork.

Synchronization of Digital Signals

To correctly understand the concepts and details of SONET, it's important to be clear about the meaning of Synchronous, Asynchronous, and Plesiochronous.

In a set of Synchronous signals, the digital transitions in the signals occur at exactly the same rate. There may, however, be a phase difference between the transitions of the two signals, and this would lie within specified limits. These phase differences may be due to propagation time delays or jitter introduced into the transmission network. In a synchronous network, all the clocks are traceable to one Stratum 1 Primary Reference Clock (PRC). The accuracy of the PRC is better than ± 1 in 10^{11} and is derived from a cesium atomic standard.

If two digital signals are Plesiochronous, their transitions occur at "almost" the same rate, with any variation being constrained within tight limits. For example, if two networks need to interwork, their clocks may be derived from two different PRCs. Although these clocks are extremely accurate, there is a difference between one clock and the other. This is known as a plesiochronous difference.

In the case of Asynchronous signals, the transitions of the signals do not necessarily occur at the same nominal rate. Asynchronous, in this case, means that the difference between two clocks is much greater than a plesiochronous difference. For example, if two clocks are derived from free-running quartz oscillators, they could be described as asynchronous.

Basic SONET Signal

SONET defines a technology for carrying many signals of different capacities through a synchronous, flexible, optical hierarchy. This is accomplished by means of a byte-interleaved multiplexing scheme. Byte-interleaving simplifies multiplexing, and offers end-to-end network management.

The first step in the SONET multiplexing process involves the generation of the lowest level or base signal. In SONET, this base signal is referred to as Synchronous Transport Signal level-1, or simply STS-1, which operates at 51.84 Mb/s. Higher-level signals are integer multiples of STS-1, creating the family of STS-N signals in Table 1. An STS-N signal is composed of N byte-interleaved STS-1 signals. This table also includes the optical counterpart for each STS-N signal, designated OC-N (Optical Carrier level-N).

Synchronous and Non-synchronous line rates and the relationships between each are shown in Tables 1 and 2.

Table 1. SONET Hierarchy

Signal	Bit Rate	Capacity
STS-1, OC-1	51.840 Mb/s	28 DS1s or 1 DS3
STS-3, OC-3	155.520 Mb/s	84 DS1s or 3 DS3s
STS-12, OC-12	622.080 Mb/s	336 DS1s or 12 DS3s
STS-48, OC-48	2488.320 Mb/s	1344 DS1s or 48 DS3s
STS-192, OC-192	9953.280 Mb/s	5376 DS1s or 192 DS3s
STS-768, OC-768	39813.12 Mb/s	21504 DS1s or 768 DS3s

STS = Synchronous Transport Signal
OC = Optical Carrier

Table 2. Non-Synchronous Hierarchy

Signal	Bit Rate	Channels
DS0	64 kb/s	1 DS0
DS1	1.544 Mb/s	24 DS0s
DS2	6.312 Mb/s	96 DS0s
DS3	44.736 Mb/s	28 DS1s

Why Synchronize?

Synchronous versus Asynchronous

Traditionally, transmission systems have been asynchronous, with each terminal in the network running on its own clock. In digital transmission, "clocking" is one of the most important considerations. Clocking means using a series of repetitive pulses to keep the bit rate of data constant and to indicate where the ones and zeroes are located in a data stream.

Since these clocks are totally free-running and not synchronized, large variations occur in the clock rate and thus the signal bit rate. For example, a DS3 signal specified at 44.736 Mb/s + 20 ppm (parts per million) can produce a variation of up to 1789 bps between one incoming DS3 and another.

Asynchronous multiplexing uses multiple stages. Signals such as asynchronous DS1s are multiplexed, extra bits are added (bit-stuffing) to account for the variations of each individual stream, and are combined with other bits (framing bits) to form a DS2 stream. Bit-stuffing is used again to multiplex up to DS3. DS3s are multiplexed up to higher rates in the same manner. At the higher asynchronous rate, they cannot be accessed without demultiplexing.

In a synchronous system, such as SONET, the average frequency of all clocks in the system will be the same (synchronous) or nearly the same (plesiochronous). Every clock can be traced back to a highly stable reference supply. Thus, the STS-1 rate remains at a nominal 51.84 Mb/s, allowing many synchronous STS-1 signals to be stacked together when multiplexed without any bit-stuffing. Thus, the STS-1s are easily accessed at a higher STS-N rate.

Low-speed synchronous virtual tributary (VT) signals are also simple to interleave and transport at higher rates. At low speeds, DS1s are transported by synchronous VT-1.5 signals at a constant rate of 1.728 Mb/s. Single-step multiplexing up to STS-1 requires no bit stuffing and VTs are easily accessed.

Pointers accommodate differences in the reference source frequencies and phase wander, and prevent frequency differences during synchronization failures.

Synchronization Hierarchy

Digital switches and digital cross-connect systems are commonly employed in the digital network synchronization hierarchy. The network is organized with a master-slave relationship with clocks of the higher-level nodes feeding timing signals to clocks of the lower-level nodes. All nodes can be traced to a primary reference source, a Stratum 1 atomic clock with extremely high stability and accuracy. Less stable clocks are adequate to support the lower nodes.

Synchronizing SONET

The internal clock of a SONET terminal may derive its timing signal from a Building Integrated Timing Supply (BITS) used by switching systems and other equipment. Thus, this terminal will serve as a master for other SONET nodes, providing timing on its outgoing OC-N signal. Other SONET nodes will operate in a slave mode called "loop timing" with their internal clocks timed by the incoming OC-N signal. Current standards specify that a SONET network must be able to derive its timing from a Stratum 3 or higher clock.

Frame Format Structure

SONET uses a basic transmission rate of STS-1 – equivalent to 51.84 Mb/s. Higher-level signals are integer multiples of the base rate. For example, STS-3 is three times the rate of STS-1 (3 x 51.84 = 155.52 Mb/s). An STS-12 rate would be 12 x 51.84 = 622.08 Mb/s.

STS-1 Building Block

The frame format of the STS-1 signal is shown in Figure 1. In general, the frame can be divided into two main areas: Transport overhead and the Synchronous Payload Envelope (SPE).

The synchronous payload envelope can also by divided into two parts: STS path overhead and the payload. The payload is the revenue-producing traffic being transported and routed over the SONET network. Once the payload is multiplexed into the synchronous payload envelope, it can be transported and switched through SONET without having to be examined and possibly demultiplexed at intermediate nodes. Thus, SONET is said to be service-independent or transparent.
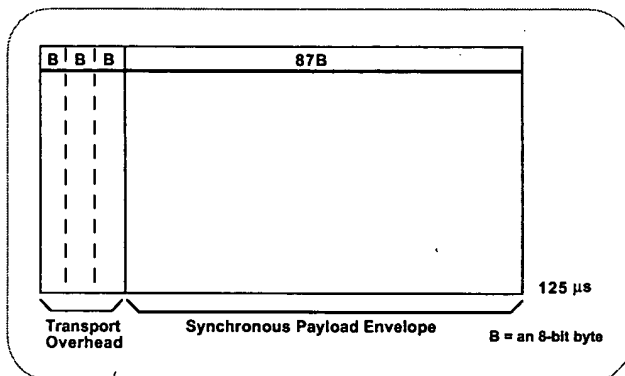
► *Figure 1. STS-1 frame format.*
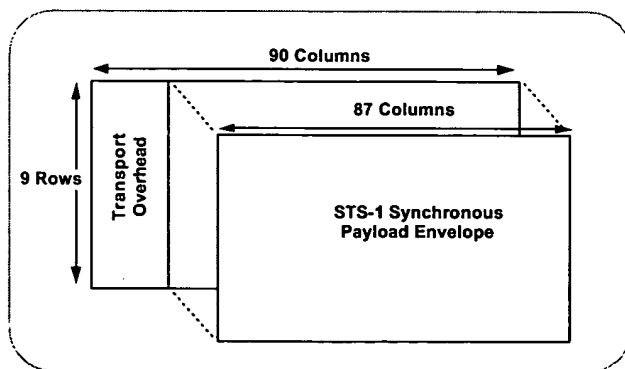
► *Figure 2. STS-1 frame elements.*

Transport Overhead is composed of section overhead and line overhead. The STS-1 path overhead is part of the synchronous payload envelope.

The STS-1 payload has the capacity to transport up to:

► 28 DS1s

► 1 DS3

► 21 2.048 Mb/s signals or combinations of the above.

STS-1 Frame Structure

STS-1 is a specific sequence of 810 bytes (6480 bits), which includes various overhead bytes and an envelope capacity for transporting payloads. It can be depicted as a 90 column by 9 row structure. With a frame length of 125 µs (8000 frames per second), STS-1 has a bit rate of 51.840 Mb/s. The order of transmission of bytes is row-by-row from top to bottom, left to right (most significant bit first).

As shown in Figure 1, the first three columns of the STS-1 frame are for the Transport Overhead. The three columns each contain nine bytes. Of these, nine bytes are overhead for the Section layer (for example, Section Overhead), and 18 bytes are overhead for the Line layer (for example, Line Overhead). The remaining 87 columns constitute the STS-1 Envelope Capacity (payload and path overhead).

As stated before, the basic signal of SONET is the Synchronous Transport Signal level 1, or STS-1. The STS frame format is composed of 9 rows of 90 columns of 8-bit bytes, or 810 bytes. The byte transmission order is row-by-row, left to right. At a rate of 8000 frames per second, that works out to a rate of 51.840 Mb/s, as the following equation demonstrates:

$$9 \times 90 \text{ bytes/frame} \times 8 \text{ bits/byte} \times 8000 \text{ frames/s} = 51,840,000 \text{ bits/s} = 51.840 \text{ Mb/s}$$

This is known as the STS-1 signal rate – the electrical rate used primarily for transport within a specific piece of hardware. The optical equivalent of STS-1 is known as OC-1, and it's used for transmission across the fiber.

SONET Telecommunications Standard

► Primer

The STS-1 frame consists of overhead, plus a Synchronous Payload Envelope (see Figure 2). The first three columns of each STS-1 frame make up the Transport Overhead, and the last 87 columns make up the SPE. SPEs can have any alignment within the frame, and this alignment is indicated by the H1 and H2 pointer bytes in the line overhead.

STS-1 Envelope Capacity and Synchronous Payload Envelope (SPE)

Figure 3 depicts the STS-1 SPE, which occupies the STS-1 Envelope Capacity. The STS-1 SPE consists of 783 bytes, and can be depicted as an 87 column by 9 row structure. Column 1 contains nine bytes, designated as the STS Path Overhead (POH). Two columns (columns 30 and 59) are not used for payload, but are designated as the "fixed stuff" columns. The 756 bytes in the remaining 84 columns are designated as the STS-1 Payload Capacity.

STS-1 SPE in Interior of STS-1 Frames

The STS-1 SPE may begin anywhere in the STS-1 Envelope Capacity (see Figure 4). Typically, it begins in one STS-1 frame and ends in the next. The STS Payload Pointer contained in the Transport Overhead designates the location of the byte where the STS-1 SPE begins.

STS POH is associated with each payload and is used to communicate various information from the point where a payload is mapped into the STS-1 SPE to where it's delivered.

STS-N Frame Structure

An STS-N is a specific sequence of N x 810 bytes. The STS-N is formed by byte-interleaving STS-1 modules (see Figure 5). The Transport Overhead of the individual STS-1 modules are frame aligned before interleaving, but the associated STS SPEs are not required to be aligned because each STS-1 has a Payload Pointer to indicate the location of the SPE (or to indicate concatenation).
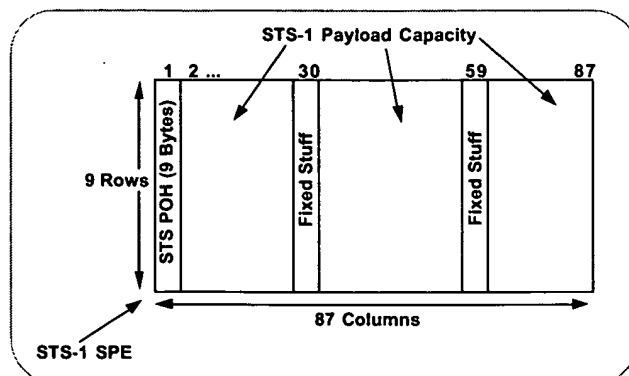
► *Figure 3. STS-1 SPE example.*
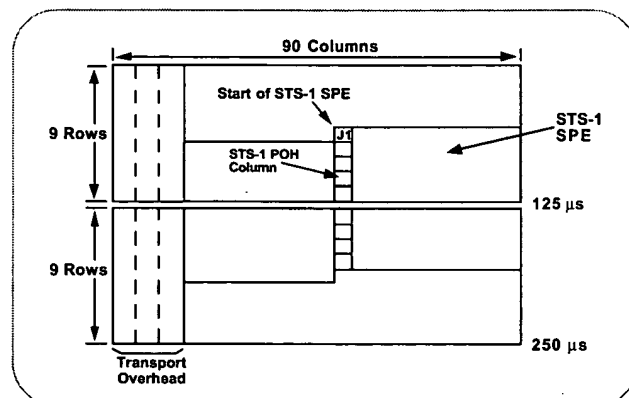
► *Figure 4. STS-1 SPE position in the STS-1 frame.*
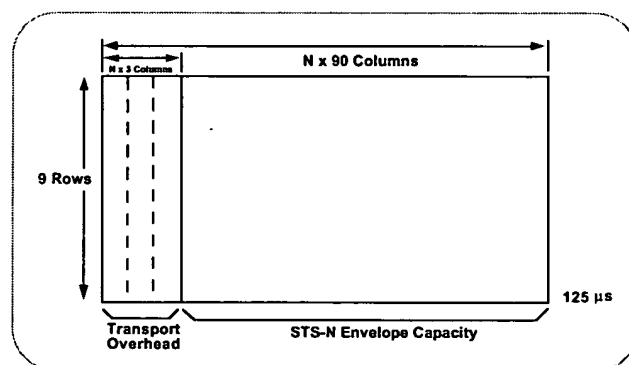
► *Figure 5. STS-N frame structure.*

Overheads

SONET provides substantial overhead information, allowing simpler multiplexing and greatly expanded OAM&P (Operations, Administration, Maintenance, and Provisioning) capabilities. The overhead information has several layers, which are shown in Figure 6. Path-level overhead is carried from end-to-end; it's added to DS1 signals when they are mapped into virtual tributaries and for STS-1 payloads that travel end-to-end. Line overhead is for the STS-N signal between STS-N multiplexers. Section overhead is used for communications between adjacent network elements, such as regenerators.

Enough information is contained in the overhead to allow the network to operate and allow OAM&P communications between an intelligent network controller and the individual nodes.

The following sections detail the different SONET overhead information:

► Section Overhead

► Line Overhead

► STS Path Overhead

► VT Path Overhead
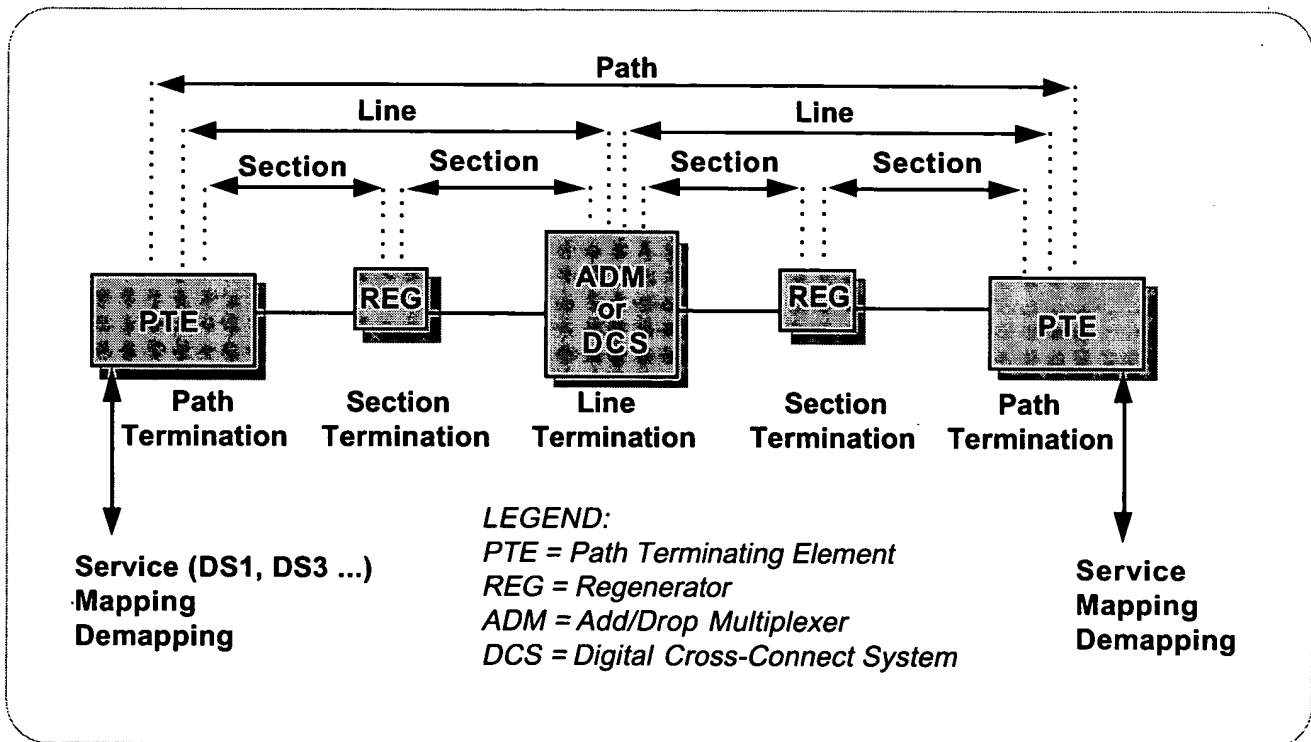
► *Figure 6. Overhead layers.*

SONET Telecommunications Standard
► Primer

Section Overhead

Section Overhead contains nine bytes of the transport overhead accessed, generated, and processed by section-terminating equipment. This overhead supports functions such as:

- ► Performance monitoring (STS-N signal)
- ► Local orderwire
- ► Data communication channels to carry information for OAM&P
- ► Framing

This might be two regenerators, line terminating equipment, and a regenerator, or two line terminating equipment. The Section Overhead is found in the first three rows of Columns 1 through 9 (see Figure 7). Table 3 shows Section Overhead byte-by-byte.

Line Overhead

Line Overhead contains 18 bytes of overhead accessed, generated, and processed by line terminating equipment. This overhead supports functions such as:

- ► Locating the SPE in the frame
- ► Multiplexing or concatenating signals
- ► Performance monitoring
- ► Automatic protection switching
- ► Line maintenance

The Line Overhead is found in Rows 4 to 9 of Columns 1 through 9 (see Figure 8). Table 4 shows Line Overhead byte-by-byte.

Table 3. Section Overhead

Byte	Description
A1 and A2	**Framing bytes** – These two bytes indicate the beginning of an STS-1 frame.
J0	**Section Trace (J0)/Section Growth (Z0)** – The byte in each of the N STS-1s in an STS-N that was formerly defined as the STS-1 ID (C1) byte has been refined either as the Section Trace byte (in the first STS-1 of the STS-N), or as a Section Growth byte (in the second through Nth STS-1s).
B1	**Section bit interleaved parity code (BIP-8) byte** – This is a parity code (even parity) used to check for transmission errors over a regenerator section. Its value is calculated over all bits of the previous STS-N frame after scrambling, then placed in the B1 byte of STS-1 before scrambling. Therefore, this byte is defined only for STS-1 number 1 of an STS-N signal.
E1	**Section orderwire byte** – This byte is allocated to be used as a local orderwire channel for voice communication between regenerators.
F1	**Section user channel byte** – This byte is set aside for users' purposes. It can be read and/or written to at each section terminating equipment in that line.
D1, D2, D3	**Section data communications channel (DCC) bytes** – These three bytes form a 192 kb/s message channel providing a message-based channel for Operations, Administration, Maintenance, and Provisioning (OAM&P) between pieces of section-terminating equipment. The channel is used from a central location for control, monitoring, administration, and other communication needs.

► *Figure 7. Section Overhead – Rows 1 to 3 of Transport Overhead.*

► *Figure 8. Line Overhead – Rows 4 to 9 of Transport Overhead.*

Table 4. Line Overhead

Byte	Description
H1, H2	**STS Payload Pointer (H1 and H2)** – Two bytes are allocated to a pointer that indicates the offset in bytes between the pointer and the first byte of the STS SPE. The pointer bytes are used in all STS-1s within an STS-N to align the STS-1 Transport Overhead in the STS-N, and to perform frequency justification. These bytes are also used to indicate concatenation, and to detect STS Path Alarm Indication Signals (AIS-P).
H3	**Pointer Action Byte (H3)** – The pointer action byte is allocated for SPE frequency justification purposes. The H3 byte is used in all STS-1s within an STS-N to carry the extra SPE byte in the event of a negative pointer adjustment. The value contained in this byte when it's not used to carry the SPE byte is undefined.
B2	**Line bit interleaved parity code (BIP-8) byte** – This parity code byte is used to determine if a transmission error has occurred over a line. It's even parity, and is calculated over all bits of the Line Overhead and STS-1 SPE of the previous STS-1 frame before scrambling. The value is placed in the B2 byte of the Line Overhead before scrambling. This byte is provided in all STS-1 signals in an STS-N signal.

| K1 and K2 | **Automatic Protection Switching (APS channel) bytes** – These two bytes are used for Protection Signaling between Line Terminating entities for bi-directional automatic protection switching and for detecting alarm indication signal (AIS-L) and Remote Defect Indication (RDI) signals. |

K1 Byte

Bits 1-4	Type of request
1111	Lock out of Protection
1110	Forced Switch
1101	SF – High Priority
1100	SF – Low Priority
1011	SD – High Priority
1010	SD – Low Priority
1001	(not used)
1000	Manual Switch
0111	(not used)
0110	Wait-to-Restore
0101	(not used)
0100	Exercise
0011	(not used)
0010	Reverse Request
0001	Do Not Revert
0000	No Request
Bits 5-8	Indicate the number of the channel requested

K2 Byte

Bits 1-4	Selects channel number
Bit 5	Indication of architecture
0	1+1
1	1:n
Bit 6-8	Mode of operation
111	AIS-L
110	RDI-L
101	Provisioned mode is bidirectional
100	Provisioned mode is unidirectional
011	Future use
010	Future use
001	Future use
000	Future use

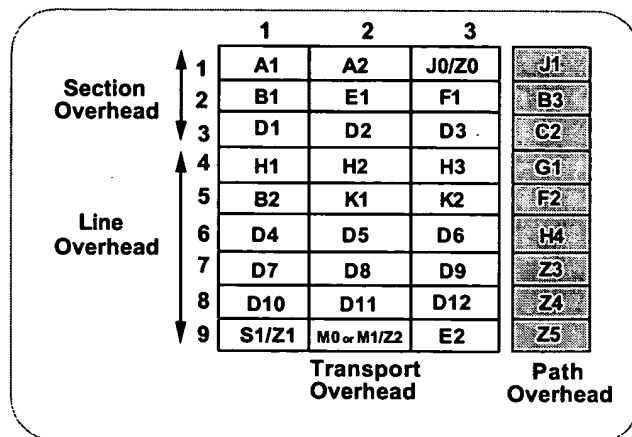| Byte	Description
D4 to D12	**Line Data Communications Channel (DCC) bytes** – These nine bytes form a 576 kb/s message channel from a central location for OAM&P information (alarms, control, maintenance, remote provisioning, monitoring, administration, and other communication needs) between line entities. A protocol analyzer is required to access the Line-DCC information.
S1	**Synchronization Status (S1)** – The S1 byte is located in the first STS-1 of an STS-N, and bits 5 through 8 of that byte are allocated to convey the synchronization status of the network element.
Z1	**Growth (Z1)** – The Z1 byte is located in the second through Nth STS-1s of an STS-N (3≤N≤48), and is allocated for future growth. Note that an OC-1 or STS-1 electrical signal does not contain a Z1 byte.
M0	**STS-1 REI-L (M0)** – The M0 byte is only defined for STS-1 in an OC-1 or STS-1 electrical signal. Bits 5 through 8 are allocated for a Line Remote Error Indication function (REI-L – formerly referred to as Line FEBE), which conveys the error count detected by an LTE (using the Line BIP-8 code) back to its peer LTE.
M1	**STS-N REI-L (M1)** – The M1 byte is located in the third STS-1 (in order of appearance in the byte-interleaved STS-N electrical or OC-N signal) in an STS-N (N≥3), and is used for a REI-L function.
Z2	**Growth (Z2)** – The Z2 byte is located in the first and second STS-1s of an STS-3, and the first, second, and fourth through Nth STS-1s of an STS-N (12≤N≤48). These bytes are allocated for future growth. Note that an OC-1 or STS-1 electrical signal does not contain a Z2 byte.
E2	**Orderwire byte** – This orderwire byte provides a 64 kb/s channel between line entities for an express orderwire. It's a voice channel for use by technicians and will be ignored as it passes through the regenerators.

STS Path Overhead

STS Path Overhead (STS POH) contains nine evenly distributed Path Overhead bytes per 125 microseconds starting at the first byte of the STS SPE. STS POH provides for communication between the point of creation of an STS SPE and its point of disassembly. This overhead supports functions such as:

► Performance monitoring of the STS SPE

► Signal label (the content of the STS SPE, including status of mapped payloads)

► Path status

► Path trace

The Path Overhead is found in Rows 1 to 9 of the first column of the STS-1 SPE (see Figure 9). Table 5 describes Path Overhead byte-by-byte.

		1	2	3	
Section Overhead	1	A1	A2	J0/Z0	J1
	2	B1	E1	F1	B3
	3	D1	D2	D3	C2
Line Overhead	4	H1	H2	H3	G1
	5	B2	K1	K2	F2
	6	D4	D5	D6	H4
	7	D7	D8	D9	Z3
	8	D10	D11	D12	Z4
	9	S1/Z1	M0 or M1/Z2	E2	Z5
		Transport Overhead			**Path Overhead**

► *Figure 9. Path Overhead – Rows 1 to 9.*

Table 5. STS Path Overhead

Byte	Description
J1	**STS path trace byte** – This user-programmable byte repetitively transmits a 64-byte, or 16-byte E.164 format string. This allows the receiving terminal in a path to verify its continued connection to the intended transmitting terminal.
B3	**STS Path Bit Interleaved Parity code (Path BIP-8) byte** – This is a parity code (even), used to determine if a transmission error has occurred over a path. Its value is calculated over all the bits of the previous synchronous payload envelope (SPE) before scrambling and placed in the B3 byte of the current frame.

C2	**STS Path signal label byte** – This byte is used to indicate the content of the STS SPE, including the status of the mapped payloads.

Bits 1-4	Bits 5-8	Status
0000	0000	Unequipped
0000	0001	Equipped
0000	0010	VT-structured STS-1 SPE
0000	0011	Locked VT mode
0000	0100	Asynchronous mapping for DS3
0001	0010	Asynchronous mapping for DS4NA
0001	0011	Mapping for ATM
0001	0100	Mapping for DQDB
0001	0101	Asynchronous mapping for FDDI

G1 — **Path status byte** – This byte is used to convey the path terminating status and performance back to the originating path terminating equipment. Therefore, the duplex path in its entirety can be monitored from either end, or from any point along the path. Bits 1 through 4 are allocated for an STS Path REI function (REI-P – formerly referred to as STS Path FEBE). Bits 5, 6, and 7 of the G1 byte are allocated for an STS Path RDI (RDI-P) signal. Bit 8 of the G1 byte is currently undefined.

	REI-P				RDI-P		undefined
1	2	3	4	5	6	7	8

Bits 1-4	STS Path REI function	
Bits 5-7	STS Path RDI function (Trigger & Interpretation)	
111	AIS-P, LOP-P	Remote defect
110	UNEQ-P, TIM-P	Remote connectivity defect
101	AIS-P, LOP-P	Remote server defect
100	AIS-P, LOP-P	Remote defect
011	No defects	No remote defect
010	PLM-P, LCD-P	Remote payload defect
001	No defects	No remote defect
000	No defects	No remote defect
Bit 8	Undefined	

F2	**Path user channel byte** – This byte is used for user communication between path elements.
H4	**Virtual Tributary (VT) multiframe indicator byte** – This byte provides a generalized multiframe indicator for payload containers. At present, it's used only for tributary unit structured payloads.

NOTE: The Path Overhead Portion of the SPE remains with the payload until it's demultiplexed.

VT Path Overhead

VT Path Overhead (VT POH) contains four evenly distributed Path Overhead bytes per VT SPE starting at the first byte of the VT SPE. VT POH provides communication between the point of creation of an VT SPE and its point of disassembly.

Four bytes (V5, J2, Z6, and Z7) are allocated for VT POH. The first byte of a VT SPE (i.e., the byte in the location pointed to by the VT Payload Pointer) is the V5 byte, while the J2, Z6, and Z7 bytes occupy the corresponding locations in the subsequent 125 microsecond frames of the VT Superframe.

The V5 byte provides the same functions for VT paths that the B3, C2, and G1 bytes provide for STS paths; namely error checking, signal label, and path status. The bit assignments for the V5 byte are described in Table 6.

Table 6. VT Path Overhead

Byte	Description
V5	**VT path overhead byte.**

BIP-2		REI-V	RFI-V	Signal Label			RDI-V
1	2	3	4	5	6	7	8

Bits 1-2	Allocated for error performance monitoring.
Bit 3	Allocated for a VT Path REI function (REI-V – formerly referred to as VT Path FEBE) to convey the VT Path terminating performance back to an originating VT PTE.
Bit 4	Allocated for a VT Path Remote Failure Indication (RFI-V) in the byte-synchronous DS1 mapping.
Bits 5-7	Allocated for a VT Path Signal Label to indicate the content of the VT SPE.
000	Unequipped
001	Equipped – non-specific payload
010	Asynchronous mapping
011	Bit synchronous mapping (no longer valid for DS1)
100	Byte synchronous mapping
101	Unassigned
110	Unassigned
111	Unassigned
Bit 8	Allocated for a VT Path Remote Defect Indication (RDI-V) signal.

Byte	Description
J2	**VT Path trace identifier** – This byte is used to support the end-to-end monitoring of a path.
Z6	The Z6 byte is known as N2 in the SDH standard and is allocated to provide a Lower-Order Tandem Connection Monitoring (LO-TCM) function.
Z7	The Z7 byte is known as K4 in the SDH standard. Bits 1-4 are allocated for APS signaling for protection at the Lower-Order path level. Bits 5-7 are used in combination with V5 bit 8 for ERDI-V. Bit 8 is reserved for future use and has no defined value.

SONET Alarm Structure

The SONET frame structure has been designed to contain a large amount of overhead information. The overhead information provides a variety of management and other functions such as:

- ▶ Error performance monitoring
- ▶ Pointer adjustment information
- ▶ Path status
- ▶ Path trace
- ▶ Section trace
- ▶ Remote defect, error, and failure indications
- ▶ Signal labels
- ▶ New data flag indications
- ▶ Data communications channels (DCC)
- ▶ Automatic Protection Switching (APS) control
- ▶ Orderwire
- ▶ Synchronization status message

Much of this overhead information is involved with alarm and in-service monitoring of the particular SONET sections.

SONET alarms are defined as follows:

Anomaly – The smallest discrepancy which can be observed between the actual and desired characteristics of an item. The occurrence of a single anomaly does not constitute an interruption in the ability to perform a required function.

Defect – The density of anomalies has reached a level where the ability to perform a required function has been interrupted. Defects are used as input for performance monitoring, the control of consequent actions, and the determination of fault cause.

Failure – The inability of a function to perform a required action persisted beyond the maximum time allocated.

Table 7 describes SONET alarm anomalies, defects, and failures.

Table 7. Anomalies, Defects, and Failures

Abbreviation	Description	Criteria
LOS	Loss of Signal	LOS is raised when the synchronous signal (STS-N) level drops below the threshold at which a BER of 1 in 103 is predicted. It could be due to a cut cable, excessive attenuation of the signal, or equipment fault. The LOS state clears when two consecutive framing patterns are received and no new LOS conditions detected.
OOF	Out of Frame Alignment	OOF state occurs when four or five consecutive SONET frames are received with invalid (errored) framing patterns (A1 and A2 bytes). The maximum time to detect OOF is 625 microseconds. OOF state clears when two consecutive SONET frames are received with valid framing patterns.
LOF	Loss of Frame Alignment	LOF state occurs when the OOF state exists for a specified time in milliseconds. The LOF state clears when an in-frame condition exists continuously for a specified time in milliseconds.
LOP	Loss of Pointer	LOP state occurs when N consecutive invalid pointers are received or "N" consecutive New Data Flags (NDF) are received (other than in a concatenation indicator), where N = 8, 9, or 10. LOP state is cleared when three equal valid pointers or three consecutive AIS indications are received. LOP can also be identified as: ▶ LOP-P (STS Path Loss of Pointer) ▶ LOP-V (VT Path Loss of Pointer)
AIS	Alarm Indication Signal	The AIS is an all-ONES characteristic or adapted information signal. It's generated to replace the normal traffic signal when it contains a defect condition in order to prevent consequential downstream failures being declared or alarms being raised. AIS can also be identified as: ▶ AIS-L (Line Alarm Indication Signal) ▶ AIS-P (STS Path Alarm Indication Signal) ▶ AIS-V (VT Path Alarm Indication Signal)
REI	Remote Error Indication	An indication returned to a transmitting node (source) that an errored block has been detected at the receiving node (sink). This indication was formerly known as Far End Block Error (FEBE). REI can also be identified as: ▶ REI-L (Line Remote Error Indication) ▶ REI-P (STS Path Remote Error Indication) ▶ REI-V (VT Path Remote Error Indication)
RDI	Remote Defect Indication	A signal returned to the transmitting Terminating Equipment upon detecting a Loss of Signal, Loss of Frame, or AIS defect. RDI was previously known as FERF. RDI can also be identified as: ▶ RDI-L (Line Remote Defect Indication) ▶ RDI-P (STS Path Remote Defect Indication) ▶ RDI-V (VT Path Remote Defect Indication)
RFI	Remote Failure Indication	A failure is a defect that persists beyond the maximum time allocated to the transmission system protection mechanisms. When this situation occurs, an RFI is sent to the far end and will initiate a protection switch if this function has been enabled. RFI can also be identified as: ▶ RFI-L (Line Remote Failure Indication) ▶ RFI-P (STS Path Remote Failure Indication) ▶ RFI-V (VT Path Remote Failure Indication)
B1 error	B1 error	Parity errors evaluated by byte B1 (BIP-8) of an STS-N are monitored. If any of the eight parity checks fail, the corresponding block is assumed to be in error.
B2 error	B2 error	Parity errors evaluated by byte B2 (BIP-24 x N) of an STS-N are monitored. If any of the N x 24 parity checks fail, the corresponding block is assumed to be in error.
B3 error	B3 error	Parity errors evaluated by byte B3 (BIP-8) of a VT-N (N = 3, 4) are monitored. If any of the eight parity checks fail, the corresponding block is assumed to be in error.
BIP-2 error	BIP-2 error	Parity errors contained in bits 1 and 2 (BIP-2: Bit Interleaved Parity-2) of byte V5 of an VT-M (M = 11, 12, 2) is monitored. If any of the two parity checks fail, the corresponding block is assumed to be in error.
LSS	Loss of Sequence Synchronization	Bit error measurements using pseudo-random sequences can only be performed if the reference sequence produced on the receiving side of the test set-up is correctly synchronized to the sequence coming from the object under test. In order to achieve compatible measurement results, it's necessary that the sequence synchronization characteristics are specified. Sequence synchronization is considered to be lost and resynchronization shall be started if: ▶ Bit error ratio is ≥0.20 during an integration interval of 1 second; or ▶ It can be unambiguously identified that the test sequence and the reference sequence are out of phase

NOTE: One method to recognize the out-of-phase condition is evaluation of the error pattern resulting from the bit-by-bit comparison. If the error pattern has the same structure as the pseudo-random test sequence, the out-of-phase condition is reached.

Pointers

SONET uses a concept called "pointers" to compensate for frequency and phase variations. Pointers allow the transparent transport of synchronous payload envelopes (either STS or VT) across plesiochronous boundaries (that is, between nodes with separate network clocks having almost the same timing). The use of pointers avoids the delays and loss of data associated with the use of large (125-microsecond frame) slip buffers for synchronization.

Pointers provide a simple means of dynamically and flexibly phase-aligning STS and VT payloads, thereby permitting ease of dropping, inserting, and cross-connecting these payloads in the network. Transmission signal wander and jitter can also be readily minimized with pointers.

Figure 10 shows an STS-1 pointer (H1 and H2 bytes) which allows the SPE to be separated from the transport overhead. The pointer is simply an offset value that points to the byte where the SPE begins. The diagram depicts the typical case of the SPE overlapping onto two STS-1 frames. If there are any frequency or phase variations between the STS-1 frame and its SPE, the pointer value will be increased or decreased accordingly to maintain synchronization.

VT Mappings

There are several options for how payloads are actually mapped into the VT. Locked-mode VTs bypass the pointers with a fixed byte-oriented mapping of limited flexibility. Floating mode mappings use the pointers to allow the payload to float within the VT payload. There are three different floating mode mappings – asynchronous, bit-synchronous, and byte-synchronous.

Concatenated Payloads

For future services, the STS-1 may not have enough capacity to carry some services. SONET offers the flexibility of concatenating STS-1s to provide the necessary bandwidth. Consult the Glossary for an explanation of concatenation. STS-1s can be concatenated up to STS-3c. Beyond STS-3, concatenation is done in multiples of STS-3c. Virtual tributaries can be concatenated up to VT-6 in increments of VT-1.5, VT-2, or VT-6.

Payload Pointers

When there's a difference in phase or frequency, the pointer value is adjusted. To accomplish this, a process known as byte stuffing is used. In other words, the SPE payload pointer indicates where in the container capacity a VT starts, and the byte stuffing process allows dynamic alignment of the SPE in case it slips in time.
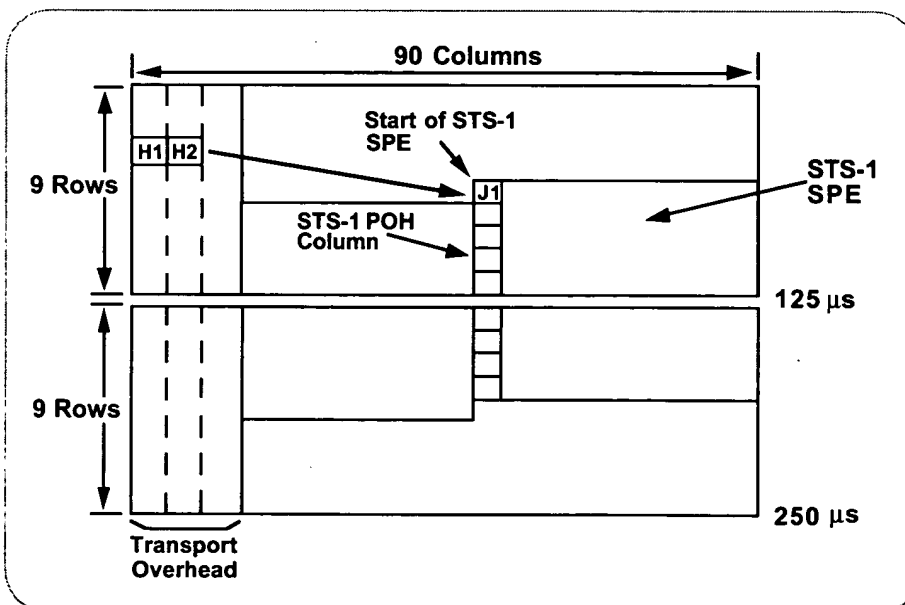

► *Figure 10. Pointer – SPE position in the STS-1 frame.*

Positive Stuffing

When the frame rate of the SPE is too slow in relation to the rate of the STS-1, bits 7, 9, 11, 13, and 15 of the pointer word are inverted in one frame, thus allowing 5-bit majority voting at the receiver. (These bits are known as the I-bits or Increment bits.) Periodically, when the SPE is about one byte off, these bits are inverted, indicating that positive stuffing must occur. An additional byte is stuffed in, allowing the alignment of the container to slip back in time. This is known as positive stuffing, and the stuff byte is made up of non-information bits. The actual positive stuff byte immediately follows the H3 byte (that is, the stuff byte is within the SPE portion). The pointer is incremented by one in the next frame, and the subsequent pointers contain the new value. Simply put, if the SPE frame is traveling more slowly than the STS-1 frame, every now and then "stuffing" an extra byte in the flow gives the SPE a one-byte delay. See Figure 11.
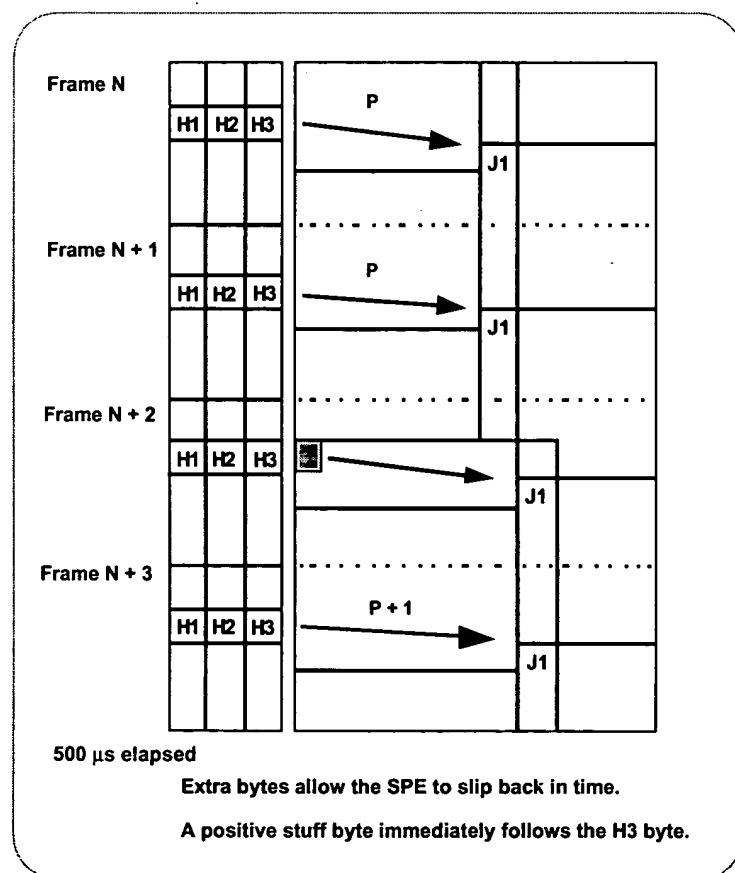
Negative Stuffing

Conversely, when the frame rate of the SPE frame is too fast in relation to the rate of the STS-1 frame, bits 8, 10, 12, 14, and 16 of the pointer word are inverted, thus allowing 5-bit majority voting at the receiver. (These bits are known as the D-bits, or Decrement bits.) Periodically, when the SPE frame is about one byte off, these bits are inverted, indicating that negative stuffing must occur. Because the alignment of the container advances in time, the envelope capacity must be moved forward. Thus, actual data is written in the H3 byte, the negative stuff opportunity (within the Overhead); this is known as negative stuffing.

The pointer is decremented by one in the next frame, and the subsequent pointers contain the new value. Simply put, if the SPE frame is traveling more quickly than the STS-1 frame, every now and then pulling an extra byte from the flow and stuffing it into the Overhead capacity (the H3 byte) gives the SPE a one-byte advance. In either case, there must be at least three frames in which the pointer remains constant before another stuffing operation (and therefore a pointer value change) can occur. See Figure 12.

Virtual Tributaries

In addition to the STS-1 base format, SONET also defines synchronous formats at sub-STS-1 levels. The STS-1 payload may be subdivided into virtual tributaries, which are synchronous signals used to transport lower-speed transmissions. The sizes of VTs are shown in Table 8.

In order to accommodate mixes of different VT types within an STS-1 SPE, the VTs are grouped together. An STS-1 SPE that is carrying Virtual Tributaries is divided into seven VT Groups, with each VT Group using 12 columns of the STS-1 SPE; note that the number of columns in each of the different VT types – 3, 4, 6, and 12 – are all factors of 12. Each VT Group can contain only one size (type) of Virtual Tributary, but within an STS-1 SPE, there can be a mix of the different VT Groups.

500 µs elapsed

Extra bytes allow the SPE to slip back in time.

A positive stuff byte immediately follows the H3 byte.

► *Figure 11. Payload pointer – positive justification.*

For example, an STS-1 SPE may contain four VT1.5 groups and three VT6 groups, for a total of seven VT Groups. Thus, an SPE can carry a mix of any of the seven groups. The groups have no overhead or pointers; they're just a way of organizing the different VTs within an STS-1 SPE.

Since each of the VT Groups is allocated 12 columns of the Synchronous Payload Envelope, a VT Group would contain one of the following combinations:
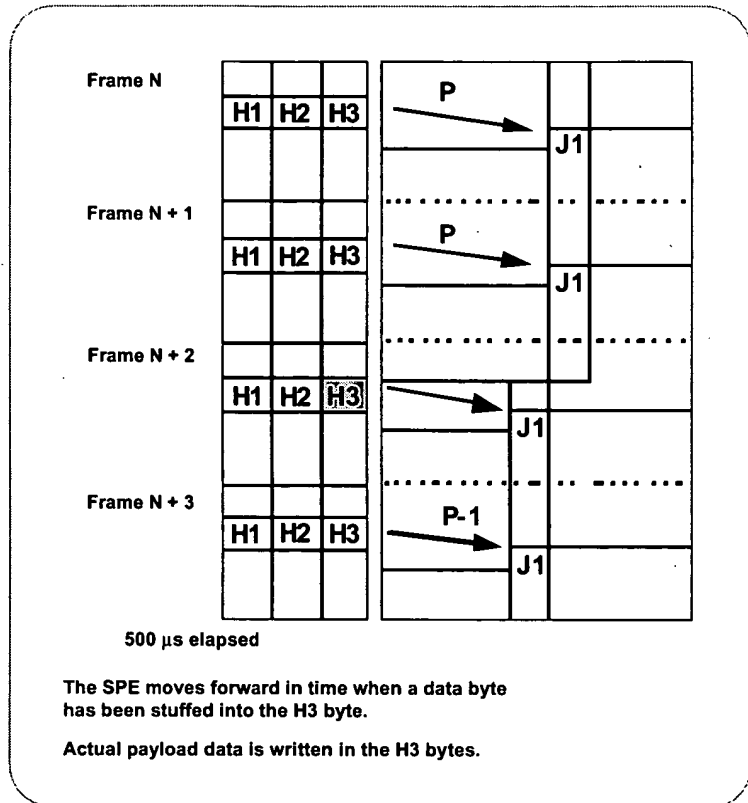
► Four VT1.5s (with 3 columns per VT1.5)

► Three VT2s (with 4 columns per VT2)

► Two VT3s (with 6 columns per VT3)

► One VT6 (with 12 columns per VT6)

The 12 columns in a VT Group are not consecutive within the SPE; they're interleaved column by column with respect to the other VT groups. As well, column 1 is used for the Path Overhead; the two columns of "fixed stuff" are assigned to columns 30 and 59.

The first VT Group, called Group 1, is found in every seventh column, starting with column 2, and skipping columns 30 and 59. That is, the 12 columns for VT Group 1 are columns 2, 9, 16, 23, 31, 38, 45, 52, 60, 67, 74, and 81.

Table 8. Virtual Tributaries (VT)

VT Type	Bit Rate	Size of VT
VT1.5	1.728 Mb/s	9 rows, 3 column
VT2	2.304 Mb/s	9 rows, 4 columns
VT3	3.456 Mb/s	9 rows, 6 columns
VT6	6.912 Mb/s	9 rows, 12 columns

500 µs elapsed

The SPE moves forward in time when a data byte has been stuffed into the H3 byte.

Actual payload data is written in the H3 bytes.

► *Figure 12. Payload pointer – negative justification.*

SONET Telecommunications Standard
► Primer

Just as the VT Group columns are not placed in consecutive columns in an STS-1 SPE, the Virtual Tributary columns within a group are not placed in consecutive columns within that group. The columns of the individual VTs within the VT Group are interleaved as well. See Figure 13.

The VT structure is designed for transport and switching of sub-STS-1 rate payloads. There are four sizes of VTs: VT1.5 (1.728 Mb/s), VT2 (2.304 Mb/s), VT3 (3.456 Mb/s), and VT6 (6.912 Mb/s). In the 87 column by 9 row structure of the STS-1 SPE, these VTs occupy columns 3, 4, 6, and 12, respectively.

To accommodate a mix of VT sizes efficiently, the VT structured STS-1 SPE is divided into seven VT groups. Each VT group occupies 12 columns of the 87 column STS-1 SPE, and may contain 4 VT1.5s, 3 VT2s, 2 VT3s, or 1 VT6. A VT group can contain only one size of VTs; however, a different VT size is allowed for each VT group in an STS-1 SPE. See Figure 14.
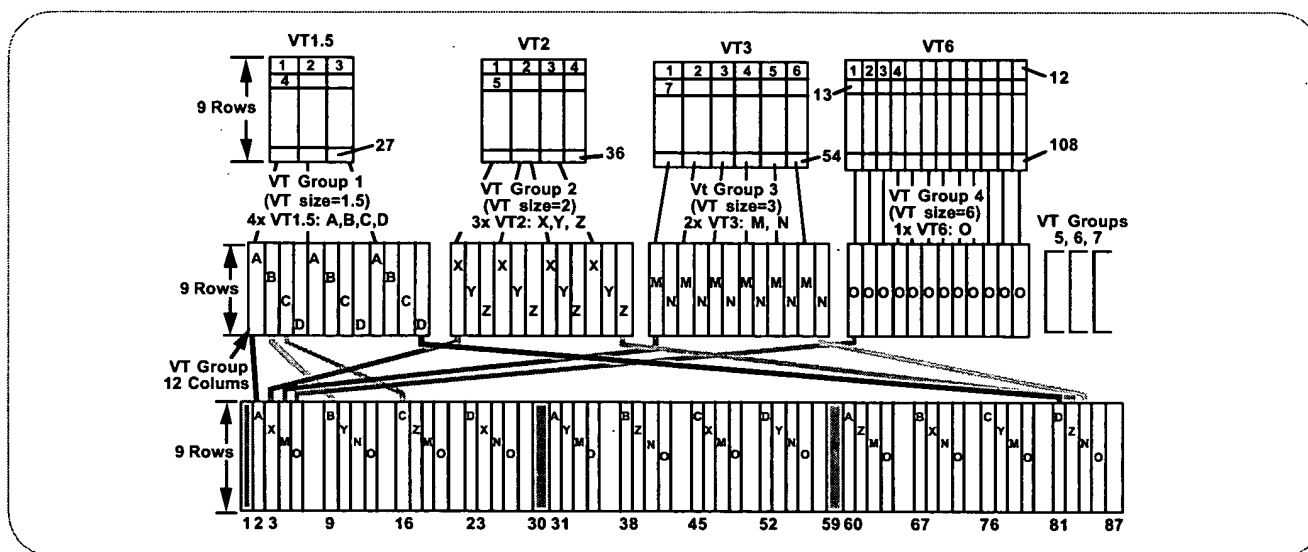
► Figure 13. SONET tributaries – VT structured STS-1 SPE.
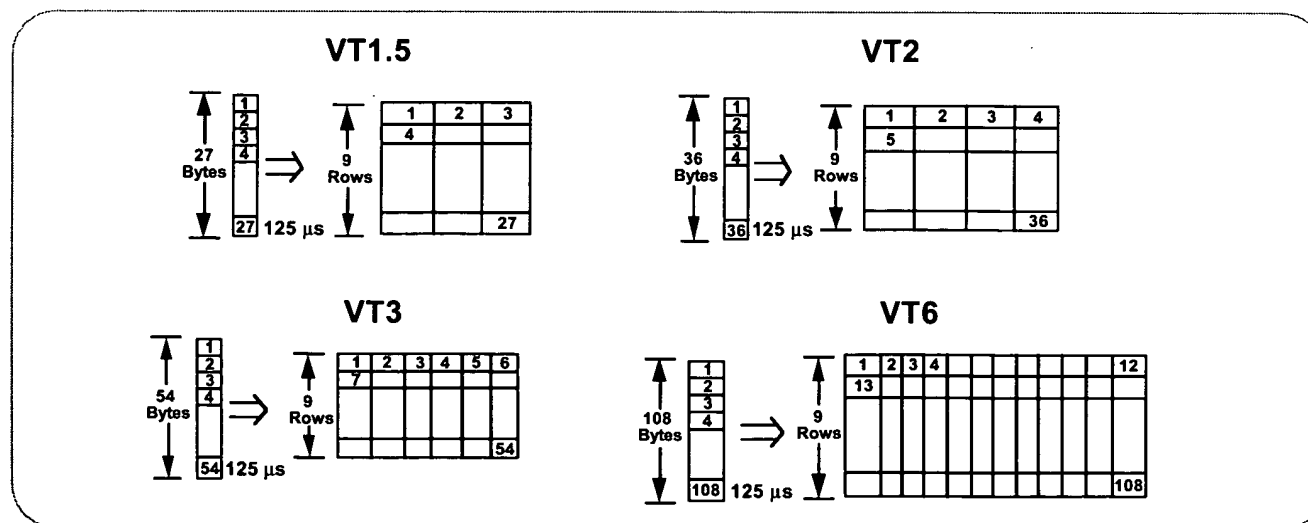
► Figure 14. VT structure, VT sizes.

STS-1 VT1.5 SPE Columns

One of the benefits of SONET is that it can carry large payloads (above 50 Mb/s). However, the existing digital hierarchy can be accommodated as well, thus protecting investments in current equipment. To achieve this capacity, the STS Synchronous Payload Envelope (SPE) can be sub-divided into smaller components or structures, known as Virtual Tributaries (VT) for the purpose of transporting and switching payloads smaller than the STS-1 rate. All services below the DS3 rate are trans-ported in the VT structure. Figure 15 shows the VT1.5 structured STS-1 SPE. Table 9 matches up the VT1.5 locations and the STS-1 SPE column numbers, per the Bellcore GR-253-CORE standard.

DS1 Visibility

Because the multiplexing is synchronous, the low-speed tributaries (input signals) can be multiplexed together but are still visible at higher rates. An individual VT containing a DS1 can be extracted without demultiplexing the entire STS-1. This improved accessibility improves switching and grooming at VT or STS levels.

In an asynchronous DS3 frame, the DS1s have gone through two levels of multiplexing (DS1 to DS2; DS2 to DS3) which include the addition of stuffing and framing bits. The DS1 signals are mixed somewhere in the information-bit fields and cannot be easily identified without completely demultiplexing the entire frame.

Different synchronizing techniques are used for multiplexing. In existing asynchronous systems, the timing for each fiber-optic transmission system termi-nal is not locked onto a common clock. Therefore, large frequency variations can occur. "Bit stuffing" is a technique used to synchro-nize the various low-speed signals to a common rate before multiplexing.

Table 9. VT1.5 Locations matched to the STS-1 SPE Column Numbers

VT Group #, VT #	Column #s
1,1	2,31,60
2,1	3,32,61
3,1	4,33,62
4,1	5,34,63
5,1	6,35,64
6,1	7,36,65
7,1	8,37,66
1,2	9,38,67
2,2	10,39,68
3,2	11,40,69
4,2	12,41,70
5,2	13,42,71
6,2	14,43,72
7,2	15,44,73
1,3	16,45,74
2,3	17,46,75
3,3	18,47,76
4,3	19,48,77
5,3	20,49,78
6,3	21,50,79
7,3	22,51,80
1,4	23,52,81
2,4	24,53,82
3,4	25,54,83
4,4	26,55,84
5,4	27,56,85
6,4	28,57,86
7,4	29,58,87

Column 1 = STS-1 POH

30 = Fixed Stuff

59 = Fixed Stuff
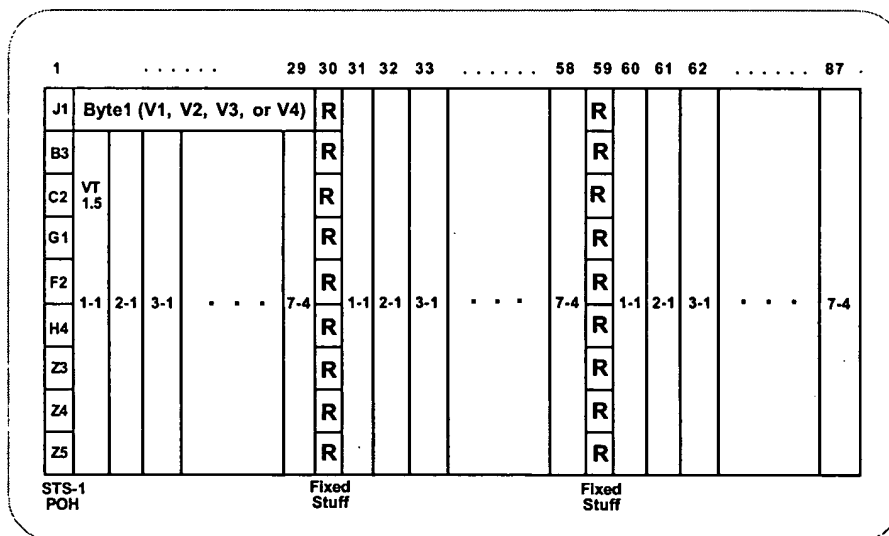

► Figure 15. STS-1 VT1.5 SPE Columns.

VT Superframe and Envelope Capacity

In addition to the division of VTs into VT groups, a 500-microsecond structure called a VT Superframe is defined for each VT. The VT Superframe contains the V1 and V2 bytes (the VT Payload Pointer), and the VT Envelope Capacity, which in turn contains the VT SPE. The VT Envelope Capacity, and therefore the size of the VT SPE, is different for each VT size. V1 is the first byte in the VT Superframe, while V2 through V4 appear as the first bytes in the following frames of the VT Superframe, regardless of the VT size. See Figure 16.
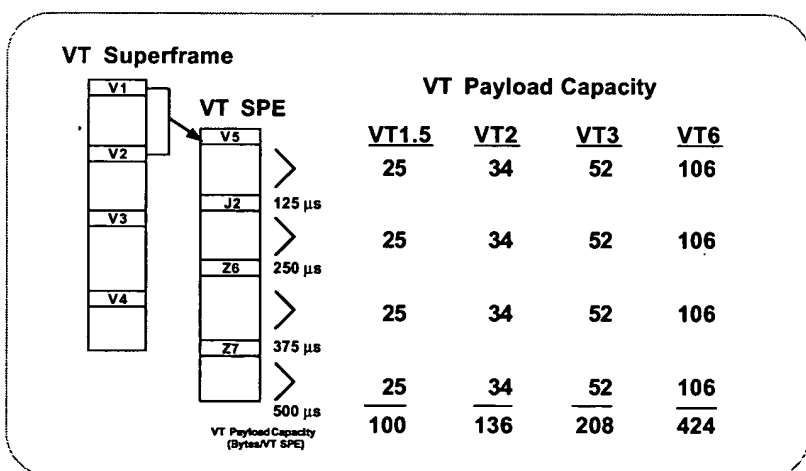
VT SPE and Payload Capacity

Four consecutive 125-microsecond frames of the VT-structured STS-1 SPE are organized into a 500-microsecond superframe, the phase of which is indicated by the H4 (Indicator) byte in the STS POH.

The VT Payload Pointer provides flexible and dynamic alignment of the VT SPE within the VT Envelope Capacity, independent of other VT SPEs. Figure 17 illustrates the VT SPEs corresponding to the four VT sizes. Each VT SPE contains four bytes of VT POH (V5, J2, Z6, and Z7), and the remaining bytes constitute the VT Payload Capacity, which is different for each VT.
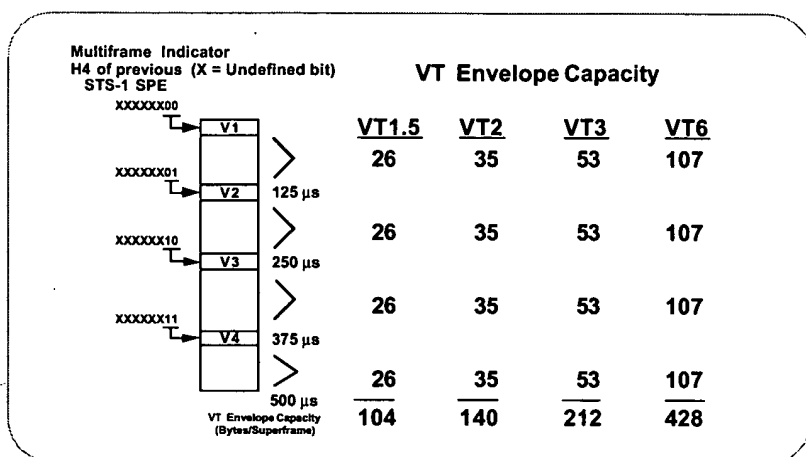
► Figure 16. VT superframe and envelope capacity.

► Figure 17. VT SPE and payload capacity.

Contact Tektronix

ASEAN Countries (65) 356-3900

Australia & New Zealand 61 (2) 9888-0100

Austria, Central Eastern Europe, Greece,
Turkey, Malta & Cyprus +43 2236 8092 0

Belgium +32 (2) 715 89 70

Brazil and South America 55 (11) 3741-8360

Canada 1 (800) 661-5625

Denmark +45 (44) 850 700

Finland +358 (9) 4783 400

France & North Africa +33 1 69 86 81 81

Germany +49 (221) 94 77 400

Hong Kong (852) 2585-6688

India (91) 80-2275577

Italy +39 (02) 25086 501

Japan (Sony/Tektronix Corporation) 81 (3) 3448-3111

Mexico, Central America & Caribbean 52 (5) 666-6333

The Netherlands +31 23 56 95555

Norway +47 22 07 07 00

People's Republic of China 86 (10) 6235 1230

Poland (48) 22 521 5340

Republic of Korea 82 (2) 528-5299

South Africa (27 11) 254-8360

Spain & Portugal +34 91 372 6000
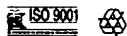
Sweden +46 8 477 65 00

Switzerland +41 (41) 729 36 40

Taiwan 886 (2) 2722-9622

United Kingdom & Eire +44 (0)1344 392000

USA 1 (800) 426-2200

For other areas, contact: Tektronix, Inc. at 1 (503) 627-1924

Tektronix
Enabling Innovation